Microsoft

SQL Server 2000 Beta 2

Relational Database Management and Analysis System

Evaluation Guide

July 2000

Abstract

This guide highlights the most important new features in Microsoft® SQL Server™ 2000. This software release includes many additional features that are not covered in this Guide.

© 2000 Microsoft Corporation. All rights reserved.

The information contained in this document represents the current view of Microsoft Corporation on the issues discussed as of the date of publication. Because Microsoft must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Microsoft, and Microsoft cannot guarantee the accuracy of any information presented after the date of publication.

This Evaluation Guide is for informational purposes only. MICROSOFT MAKES NO WARRANTIES, EXPRESS OR IMPLIED, IN THIS DOCUMENT.

Microsoft, Windows, MSDN, Windows NT, BizTalk, Win32, Visual Studio, ActiveX, Active Directory, Visual SourceSafe, Visual C++ and Visual Basic are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries

Other product or company names mentioned herein may be the trademarks of their respective owners.

Microsoft Corporation • One Microsoft Way • Redmond, WA 98052-6399 • USA

Microsoft SQL Server 2000 Evaluation Guide April 11, 2000.

CONTENTS

INTRODUCTION TO THE EVALUATION GUIDE....3
Overview 3
Audience 3

PRODUCT EVALUATION ASSISTANCE....5
Product Training 5
Additional Information 5

EVALUATING SQL SERVER 2000....7

WINDOWS DNA AND SQL SERVER 2000....9
Building Next Generation Web Applications 9
Windows 2000: The Foundation of Windows DNA 9
The Next Generation Internet Platform: Windows DNA 2000 10
Windows DNA 2000 Components 11

SQL SERVER 2000 DESIGN THEMES....13
Fully Web-Enabled 13
Highly Scalable and Reliable 14
Fastest Time-to-Market 14

FULLY WEB-ENABLED....17
Rich XML and Internet standard support 17
Easy and secure access to data via the Web 19
Secure HTTP connectivity to the database 19
Full-Text Search 21
Security 22
English Query 23
Powerful, flexible Web-based analysis 25
Data Mining and Business Internet Analytics 25
Linked Cubes and HTTP access to cubes 27
Secure application hosting 27

HIGHLY SCALABLE AND RELIABLE....29
Scalability and reliability for E-Commerce solutions 29
Distributed Partitioned Views 29
Log Shipping 30
SMP-related features 30
Scalability and reliability for Line-of-Business solutions 31
Enhanced failover clustering 31
Large memory and SMP support 31
Scalability and reliability for Data Warehousing solutions 32
Indexed views 32
Distributed Partitioned Cubes 33

Large data set analysis 33
Maximized uptime and reliability 34
More online operations 34
Differential database backups 34
Server-less snapshot backups 34

FASTEST TIME-TO-MARKET.....37
Simplified management and tuning 37
Integration with the Active Directory™ service 37
Dynamic self-management and tuning 38
Copy Database Wizard 38
Quick development, debugging and data transformation 39
Query Analyzer 39
Data Transformation Services 39
User-defined functions and programmability enhancements 40
Integrated and extensible Analysis Services 41
OLAP Actions 42
Custom rollups 42

CONCLUSION.....45

SQL Server 2000 New Features Overview.....47

Overview

This Evaluation Guide provides an overview of the design themes for Microsoft® SQL Server™ 2000 and its many new and enhanced features.

As the next major release of SQL Server, SQL Server 2000 builds upon the modern, extensible foundation of SQL Server 7.0, a critical release in Microsoft's database lineup and one in which much of the SQL Server product was both re-architected and rewritten. In the time since its release, SQL Server 7.0 has become the preferred database of many Customer Relationship Management (CRM), Business Intelligence (BI), Enterprise Resource Planning (ERP), and other line-of-business application vendors and customers because of its performance, scalability, manageability, programmability and value. Additionally, it has achieved great success as an Internet database.[1] SQL Server 2000 carries forward this tradition of excellence in database software design and engineering to become the foundation for ongoing innovation in the

[1] "Where is Web Data Stored?" Zona Research Quiniela, January 20, 2000. See: http://www.zonaresearch.com/deliverables/quinielas/indexes/2000/jan00.htm

SQL Server product line, from 64-bit and Windows®-based CE releases expected this calendar year, to “Yukon”, the next major release of SQL Server following SQL Server 2000.

Microsoft SQL Server 2000 is the complete database and analysis solution for rapidly delivering the next generation of scalable Web applications. As a core component of Windows DNA, it dramatically reduces the time required to bring e-commerce, line-of-business, and data warehousing applications to market while offering the scalability needed for the most demanding environments. SQL Server 2000 includes rich support for XML and HTTP; performance and availability features to partition load and ensure uptime; and advanced management and tuning functionality to automate routine tasks and lower total cost of ownership. Additionally, SQL Server 2000 takes full advantage of Windows 2000, including support for the Active Directory™ service, and up to 32 processors and 64 GB of RAM, with even higher limits when the 64-bit version ships later this year.

Audience

This Guide assumes that evaluators are already familiar with the basic features of relational database management systems (RDBMS), and have some experience with previous versions of SQL Server. Experience with Online Analytical Processing (OLAP) systems is also helpful.

PRODUCT EVALUATION ASSISTANCE

Product Training

Product training classes for SQL Server 2000 at Microsoft Certified Technical Education Centers will be available in the near future.

Additional Information

There are several sources of additional information for those who wish to review the new features in SQL Server 2000:

- Books Online, the documentation for SQL Server 2000, which contains both concise and in-depth sections on new features. A good starting point for evaluating SQL Server 2000 is the What's New section. Books Online is included on the SQL Server 2000 Beta 2 CD and the final release media for SQL Server 2000.
- The Microsoft SQL Server home page: www.microsoft.com/sql/ currently provides product information for SQL Server 7.0 including pricing and licensing, white papers, research reports and studies, links to third-party vendors and other resources. Additional information on SQL Server 2000 will be available after the release of SQL Server 2000 Beta 2.
- The MSDN™ Online SQL Server Developer Center: msdn.microsoft.com/sqlserver/ offers database, Web and application developers all the resources they need to build solutions with SQL Server, including online documentation, technical articles, peer advice and more.
- Microsoft TechNet: www.microsoft.com/technet provides insights and answers for IT professionals. Chats, technical forums and other resources are available for those deploying, maintaining and supporting SQL Server. These resources will be focused on SQL Server 2000 after the product's release.

.

EVALUATING SQL SERVER 2000

SQL Server 2000 is available in a variety of different editions in order to accommodate the unique performance, runtime, and price requirements of organizations and individuals.

- **SQL Server 2000 Enterprise Edition** is the complete SQL Server offering for any organization. It offers the advanced scalability and reliability features necessary for mission-critical line-of-business and Internet scenarios, including Distributed Partitioned Views, log shipping, and enhanced failover clustering. It also takes full advantage of the highest-end hardware, with support for up to 32 CPUs and 64 GB of RAM. SQL Server 2000 Enterprise Edition includes advanced analysis (OLAP) features for handling the largest cubes with many dimensions.
- **SQL Server 2000 Standard Edition** is the affordable option for small and medium-sized organizations that do not require advanced scalability and availability features or several of the more advanced analysis features of SQL Server 2000 Enterprise Edition. Standard Edition can be used on symmetric multiprocessing systems with up to 4 CPUs and 2 GB of RAM.
- **SQL Server 2000 Personal Edition** includes a full set of management tools and most of the functionality of Standard Edition, but is optimized for personal use. In addition to running on Microsoft's server operating systems, Personal Edition runs on non-server operating systems including Windows 2000 Professional, Windows NT® Workstation 4.0 and Windows 98. Dual processor systems are supported. While this edition supports databases of any size, its performance is optimized for single users and small workgroups and degrades with workloads generated by more than 5 concurrent users.
- **SQL Server 2000 Developer Edition** is the SQL Server offering that allows developers to build any type of application on top of SQL Server. It includes all of the functionality of Enterprise Edition but with a special development and test end-user license agreement (EULA) that prohibits production deployment (see the SQL Server 2000 Developer Edition EULA for complete details).
- **SQL Server 2000 Enterprise Evaluation Edition** is a complete version of Enterprise Edition. It is time limited and has other restrictions that affect its use and deployment (see the Evaluation Edition EULA for complete details).
- **SQL Server 2000 Desktop Engine (MSDE)** offers the basic database engine features of SQL Server 2000. It does not include a user interface, management tools, analysis capabilities, merge replication support, client access licenses, developer libraries, or Books Online. It also limits database size and user workload. It

has the smallest footprint of any edition of SQL Server 2000 and is thus an ideal embedded or offline data store.

- **SQL Server 2000 Windows CE Edition** is a version of SQL Server 2000 for devices and appliances running Windows CE. It is programmatically compatible with the other editions of SQL Server 2000 so developers can leverage their existing skills and applications to extend the power of a relational data store to solutions running on new classes of devices. SQL Server 2000 Windows CE Edition will be available in the second half of calendar year 2000.

WINDOWS DNA AND SQL SERVER 2000

Building Next Generation Web Applications

As the complete database and analysis solution for rapidly delivering the next generation of scalable Web applications, SQL Server 2000 is an integral piece of the Internet platform required by developers and IT staff to build and deploy solutions that integrate their customers, partners, business processes, and applications: Windows DNA. Windows DNA is the comprehensive, integrated platform for building and deploying applications for the Business Internet. From high traffic e-commerce Web sites to enterprise supply chain integration, Windows DNA provides the fastest route for Web solution development.

Windows DNA includes web application services, RAD development tools, and scalable server applications for creating web-based solutions, with Windows 2000 as the foundation.

Windows 2000: The Foundation of Windows DNA

While many vendors sell “application servers” separately, Microsoft has delivered a comprehensive set of web application services fully integrated into the Windows operating system. In other words, the Web application server for Windows DNA is Windows 2000. The following features and technologies, integrated in all editions of Windows 2000, ensure that Windows 2000 is the best foundation for building web-based solutions:

- **Internet Information Services 5.0** (IIS) – Integrated web services that enable users to easily host and manage Web sites to share information, create Web-based business applications, and extend file, print, media and communication services to the Web.
- **Active Server Pages** (ASP) – A high performance web server-scripting environment ideal for generating and presenting dynamic web content.
- **COM+ Component Services** (formerly known as Microsoft Transaction Server, or MTS) – Provides easily accessible component services for transactions, thread pooling, object pooling, just-in-time object activation, remote object invocation and more from a wide array of development languages.
- **Distributed Transactions** – Transaction support for updating two or more networked computer systems in a coordinated fashion.
- **Microsoft Message Queuing** (MSMQ) – Enables guaranteed

communication across networks and systems regardless of the current state of the communicating applications and systems.

- **Role-Based Security** – Enables developers to define security easily at the method, interface, component or package level.
- **Network Load Balancing**[2] – Provides scalability and high availability by balancing incoming IP traffic among multi-node clusters.
- **Active Directory Services** – Centralizes administrative information about users, applications and devices.
- **High-Performance XML Support (Microsoft XML Parser)** – Facilitates the creation of applications that that exchange XML-formatted data with Microsoft Internet Explorer and XML-enabled server products.
- **ActiveX® Data Objects –** A single interface to all types of data (relational and non-relational), including new support for conversion of database recordsets to XML.

These integrated web application services free developers from the burden of building or assembling the required infrastructure for distributed applications and enable them to focus on delivering business solutions. To build Windows DNA solutions—which are characterized by SQL Server on the data tier, business logic on the middle tier and either rich client or web technologies on the user tier—developers use the industry's leading development system, Microsoft Visual Studio®. With Windows DNA, developers can evolve their traditional programming skills to develop next generation Web applications, including stateless, transactional, middle-tier components and loosely coupled, message-based systems.

The Next Generation Internet Platform: Windows DNA 2000

Windows DNA 2000 is the next wave of software building blocks for the Windows DNA platform, and includes the following server products: SQL Server 2000, Application Center 2000, Commerce Server 2000, BizTalk Server 2000, and Host Integration Server 2000. Windows DNA 2000 makes it possible for organizations to build the secure, reliable, highly scalable commerce infrastructure required for business-to-consumer, business-to-business and hosted Web applications. The Windows DNA 2000 server products offer:

- **Pervasive XML and Internet Standards Support** – The Windows DNA 2000 server products use standards like XML, HTML, HTTP, FTP

[2] Microsoft® Windows® 2000 Advanced Server and Microsoft® Windows® 2000 Datacenter Server only.

and XPath to effectively exchange data across Internet and intranets, and through firewalls.

- **Software Scale Out and Scale Up** – Server "farms" or clusters divide the load of an application or system across multiple off-the-shelf PC servers to ensure availability and “just-in-time” scaling. Windows DNA 2000 extends software scale out beyond the user and middle tiers to the data tier with new multi-server partitioning support in SQL Server 2000. At the same time, applications can be run on the increasingly powerful Windows 2000 multi-processor systems developed by leading vendors that enable greater processing power on a single system, all the way up to 32-way systems supported by Windows 2000 Datacenter Server.
- **Reliability** – Downtime is inherently reduced with a software scale out approach because most, if not all, single points of failure are removed. Should a server go down, other servers pick up the load dynamically, enabling the application to continue servicing clients. Furthermore, integrated tools simplify management of all components and services across the platform, reducing management downtime and maximizing availability.
- **Interoperability** – Integrating heterogeneous systems and applications to connect customers, partners and suppliers is a core feature of the Windows DNA 2000 server products. Windows DNA 2000 allows developers to easily and quickly build solutions that integrate with all heterogeneous environments.
- **Fast Time-to-Market** – Through deep integration of all of the products and services making up the Windows DNA 2000 platform and with world-class development tools support, developers need only focus on implementing business logic when building an application. The overhead associated with rolling your own system services and infrastructure has been alleviated by including all of these technologies as standard components of the platform. This, along with a fundamental commitment to ease of use, allows developer to more quickly develop and deploy solutions than on any other platform.
- **Ease of Deployment, Administration, and Management** – In the past, building multi-tier applications and deploying and managing those applications in clustered environments has been exceedingly complex and unapproachable. A fundamental design philosophy at Microsoft is to deliver state-of-the-art technology in such a way that it is easier for developers and managers to use. Windows DNA 2000 is designed to make it easy for developers to build reliable applications that scale out on the web tier, application tier and data tier – and are inherently more manageable than applications on any previous

platform.

- **Full exploitation of Windows 2000** – Windows 2000 is the cornerstone of the Windows DNA 2000 platform and includes an integrated, enterprise-class application server. Furthermore, Windows 2000 provides an infrastructure that leverages state-of-the art hardware technologies such as Storage Area Networks (SANs), large memory and SMP architectures with up to 32 processors. Building on this platform, the rest of the Windows DNA 2000 platform leverages these services and capabilities for specific functions such as data management and XML support.

Windows DNA 2000 Components

In addition to Windows 2000 and SQL Server 2000, Windows DNA 2000 contains the following server products:

- **Microsoft Application Center 2000** – This new product simplifies the deployment and management of Windows DNA solutions within "farms" of servers. Application Center 2000 makes it easy to configure and manage high-availability server arrays.

  Just as SQL Server 2000 enables software scale out on the data tier, Application Center 2000 enables scale out for middle tier Web applications.

- **Microsoft Commerce Server 2000** – Commerce Server 2000 is the fastest way to build an effective online business. It provides all of the personalization, user and product management, marketing, closed loop analysis, and electronic ordering infrastructure necessary for both business-to-business and business-to-consumer e-commerce. Commerce Server 2000 is designed for quickly building tailored, scalable e-commerce solutions that optimize the user experience and provide business managers with real-time analysis and control of their online business.

  SQL Server 2000 provides consistent, easy-to-manage data modeling and administration tools for sites built on Commerce Server 2000. These services include user profiling, product catalogs and Business Internet Analytics, the analysis of customer Web click-stream data to make predictions about customer behavior and drive personalization.

- **Microsoft® BizTalk™ Server 2000** – A new XML-based business-to-business commerce solution, BizTalk Server 2000 provides business process integration within the enterprise and with trading partners across the Internet through the exchange of XML-formatted business documents.

SQL SERVER 2000 DESIGN THEMES

SQL Server provides high performance storage and easy-to-schedule transformation capabilities for data from BizTalk Server. SQL Server 2000 and BizTalk Server 2000 support the same XML Data Reduced schema. This allows documents to be transmitted directly from SQL Server to BizTalk Server and vice versa.

- **Microsoft Host Integration Server 2000** – This product provides network, data, and application integration with a variety of legacy hosts.

  SQL Server 2000 allows developers using Host Integration Server 2000 to store, retrieve, and transform data from legacy systems, as well as enabling bi-directional heterogeneous replication between those legacy systems and SQL Server 2000.

Microsoft makes significant investments to evaluate customer requirements and deliver products that meet those needs. From quantitative research (market sizing and tracking studies as well as in-depth user research) to qualitative investigation (focus groups and intensive feedback from customers and partners), the SQL Server team is committed to understanding existing and prospective customers. Driven by customer requests, cutting-edge research and strategic direction, SQL Server 2000 builds off of the modern foundation delivered in SQL Server 7.0. SQL Server 2000 is the complete database and analysis offering for rapidly delivering the next generation of scalable of e-commerce, line of business and data warehousing solutions. It includes features and technologies that make it:

- Fully Web-Enabled
- Highly Scalable and Reliable
- Deliver Fastest Time-to-Market

In addition to focusing on these core design themes, the SQL Server development team—at every level—has emphasized quality. In an effort to ensure that the quality of the product will exceed customer expectations, an Early Adopter Program was instituted as early builds of SQL Server 2000 became available. Early adopter partners and customers from a variety of industries have been installing and testing SQL Server 2000 for many months, even in production environments. The SQL Server team is confident that with its focus on development quality and extensive internal and external testing that SQL Server 2000 is the highest quality release of SQL Server to date.

Fully Web-Enabled

Even today, before the launch of SQL Server 2000, SQL Server is already the most popular database for storing data on the Web.[3] With SQL Server 2000, Microsoft answers customer demand for additional Web and e-commerce functionality beyond that in SQL Server 7.0. The SQL Server team has focused on:

- **Rich XML and Internet standard support.** Enable the generation and consumption of XML data in SQL Server. Allow for inserts, updates and deletes to be performed with XML. Provide a simple XML model (no complex programming) for both database and Web developers that allows them to use their current respective development skills while leveraging the power of XML. Support W3C and emerging standards.
- **Easy and secure access to data via the Web.** Allow access and querying from URLs via HTTP. Extend high performance full-text search to formatted documents through an extensible filter mechanism. Simplify development and deployment of English Query (natural language query) Web solutions.
- **Powerful, flexible Web-based analysis.** Enable existing OLAP cubes to be linked and analyzed, even over the Web. Deliver data mining tools and algorithms for Web data analysis. Integrate with Commerce Server 2000 to support complete clickstream and Web data analysis (also known as Business Internet Analytics; for details, see: www.microsoft.com/sql/bizsol/BIA.htm). Simplify DISTINCT COUNTs.
- **Secure application hosting.** Provide multi-instance support to allow multiple, isolated installations of SQL Server on a single machine. Offer cloning and distribution of disk images for efficient and consistent database creation.

Highly Scalable and Reliable

Just as SQL Server 7.0 has proven to be a popular Web and e-commerce database today, it has achieved world-record scalability in a variety of scenarios and delivered companies like barnesandnoble.com with the highest levels of availability.[4] In fact, a number of popular, high-volume sites are running on SQL Server 7.0 today; Dell.com, Nasdaq.com, Buy.com and Quote.com are just a few examples. SQL Server 2000 delivers a broad range of scalability and reliability features:

---

[3] "Where is Web Data Stored?" Zona Research Quiniela, January 20, 2000. See: http://www.zonaresearch.com/deliverables/quinielas/indexes/2000/jan00.htm

[4] barnesandnoble.com, running on SQL Server, achieved the highest average availability (98.55%) of major surveyed e-commerce sites during the early holiday shopping period according to Keynote Systems (as presented in PC Week Online, December 3, 1999).

- **Scalability for E-Commerce solutions.** Allow workload partitioning across multiple installations of SQL Server 2000. Provide automatic query optimization and support for queries across multiple servers in a SQL Server distributed database configuration. Take full advantage of Symmetric Multiprocessing (SMP) hardware.
- **Scalability for Line-of-Business (LOB) solutions.** Increase SMP support to take full advantage of new Windows 2000 capacity. Support direct access to high performance server-to-server interconnects. Partner with key ISVs to ensure high application performance on top of SQL Server 2000. Support up to 32 CPUs and 64 GB of RAM.
- **Scalability for Data Warehousing solutions.** Improve performance in the most complex reporting scenarios. Support the analysis of extremely large data sets with limited reprocessing. Allow dimensions with tens of millions of members. Scale out cubes through transparent remote partitioning to server farms.
- **Maximized uptime and reliability.** Simplify failover cluster configuration and management. Allow database to remain online during most operations. Enable differential and "server-less" snapshot backups. Integrate log shipping configuration and management.

Fastest Time-to-Market

SQL Server has long been considered the fastest way to build, deploy and manage e-commerce, line of business, and data warehousing solutions. Research studies by independent firms have demonstrated not only that SQL Server is easier to use than its primary competitor, but also that it demonstrates significantly lower total cost of ownership.[5] SQL Server 2000 includes innovations that guarantee it will continue to afford users with the greatest ease-of-use and the fastest time-to-market for their applications:

- **Integrated and extensible Analysis Services.** Provide the most complete, integrated and Web-enabled Analysis (OLAP) offering, including data mining capabilities. Allow analysis results to drive applications automatically. Provide custom roll-up capability.
- **Simplified management and tuning.** Enable management of databases centrally from the Windows 2000 Active Directory™ service. Automate management and tuning activities. Simplify movement of databases between instances and machines.

[5] Aberdeen Group found that SQL Server 7.0 offered a real world solution with 3.7x times lower TCO than the primary competition (see: http://www.microsoft.com/sql/productinfo/Aberdeen.htm). American Institutes for Research found that SQL Server 7.0 is three times easier to use than its primary competition (see: http://microsoft.com/presspass/features/1999/09-29compass2.htm).

- **Enhanced and simplified T-SQL development and debugging.** Deliver an integrated T-SQL debugger. Simplify development with T-SQL templates. Allow server-side statement tracing and client-side statistics.
- **Flexible and extensible Data Transformation Services (DTS).** Provide integration with MSMQ and enable access to data over the Internet via FTP. Deliver enhanced OLE DB support and improved error handling/recoverability.

FULLY WEB-ENABLED

E-commerce solutions put heavy demands on software, and the database is no exception. High transactional loads are the norm as users from all over the world access sites around the clock. Programmatic flexibility is critical to ensure interoperability with existing systems and the ability to respond quickly to changing business requirements. SQL Server 2000 provides extensive database programming capabilities built on Web standards while at the same time delivering the scalability and availability required for e-commerce. This combined with superior time-to-market and simplified management makes SQL Server 2000 the best database for e-commerce.

SQL Server 2000 offers features that go beyond what most e-commerce companies need today to provide a range of functionality that helps companies differentiate themselves from their competition now and well into the future. Today, e-commerce companies are realizing the need to derive further value from the massive amounts of data that are generated by users accessing their sites. The fully Web-enabled analysis capabilities of SQL Server 2000 allow companies to leverage data collected through user clickstreams to gain a deep understanding of customer behavior. Additionally, while many companies are focused on getting their e-commerce solutions to market quickly, they often don't want to deal with the day-to-day management of the software that runs their business. SQL Server 2000 is ready-to-run in a hosted scenario where a company has decided its database is best managed and run by another company, an Application Service Provider (ASP).

Rich XML and Internet standard support

In addition to being the standard technology for interchanging data on the Web, XML is quickly becoming the preferred technology for integrating e-commerce systems. Companies building business-to-consumer, business-to-business and extranet Web solutions look for XML support to ease the pain of integrating backend systems and transferring data through firewalls. XML also acts as a lingua franca that enables vertical applications from different vendors running on different platforms to communicate—including ERP and CRM software. While many companies have looked to middle-tier XML solutions to solve their data communication problems, developers have also realized the value of high-speed storage and generation capabilities for XML documents and data. SQL Server 2000 offers integrated, best-of-breed XML support that is flexible, high-performance and easy to use for both Web developers and database programmers.

The rich XML functionality in SQL Server 2000 shields Web developers from the intricacies of relational database programming; instead, they use technologies like XPath, URL queries and XML updategrams.

Similarly, database developers are not required to learn an object-oriented language or understand all the facets of XML. They can accommodate the core scenario of providing XML access to an existing relational database with the FOR XML clause that returns XML data from a SELECT statement and the OPENXML T/SQL keyword. OpenXML provides a relational view on XML data that, with T-SQL, can be used to query data from XML, join XML data with existing relational tables, and update the database.

SQL Server enables queries to return data as XML (rather than a standard rowset) from a standard SELECT statement via the FOR XML clause. The shape of the returned XML document is determined by the XML mode specified immediately after the FOR XML clause. The following XML modes can be specified:

- RAW: The RAW mode takes the query result and transforms each row in the result set into an XML element with a generic identifier row as the element tag.
- AUTO: The AUTO mode returns query results in a simple, nested XML tree. Each table in the FROM clause for which at least one column is listed in the SELECT clause is represented as an XML element. The columns listed in the SELECT clause are mapped to the appropriate element's attributes (alternatively, the ELEMENTS directive allows columns to be represented as sub-elements instead of attributes). The hierarchy (nesting of the elements) in the result set is determined based on the order of tables identified by the columns specified in the SELECT clause. The leftmost table will be the top element. The second leftmost table (identified by columns in the SELECT statement) will be nested within the top element etc.
- EXPLICIT: With EXPLICIT mode, developers can explicitly define the shape of the resulting XML tree. However, this requires that the queries be written in a specific way so that additional information about the desired nesting is specified explicitly.

Besides retrieving data as XML, it is important to be able to efficiently store data as XML, maintaining the relationships and hierarchy of data while taking full advantage of the speed offered by a high performance database like SQL Server. SQL Server 2000 can provide an XML View of relational data as well as map XML data into relational tables. SQL Server 2000 introduces annotations (extensions) to the XDR (XML-Data Reduced) schema language. These annotations are used to specify a bi-directional XML-to-relational mapping. Elements and attributes in an XML document are mapped to tables and columns in the databases. By default, an element name in an annotated schema maps to a table (or view) name in the specified database and the attribute name maps to the column name. These annotations can also be used to specify the XML hierarchical relationships. The annotations for relationship and field

mapping are:

- sql:relation: used to specify the name of a table or view when that name does not match the associated field in annotated schema.
- sql:field: specifies the mapping between <element> or <attribute> in an annotated schema to a column in a database.
- sql:overflow-field: When records are inserted in the database from an XML document, all the unconsumed data from the source XML document can be stored in a column. In retrieving data from the database using annotated schemas, the sql:overflow-field attribute can be specified to identify the column in the table in which the overflow data is stored.
- sql:relationship: is a tag is introduced to nest the annotated schema elements hierarchically based on the primary key/foreign key relationship among the underlying tables.

Whereas XML Views enable relational tables to be accessed as if they were XML documents, OpenXML allows XML documents to be addressed with relational syntax (SQL). OpenXML is a Transact-SQL (T-SQL) keyword that provides an updateable rowset over in-memory XML documents. The records in the rowset can be stored in database tables (similar to the rowsets provided by tables and views). OpenXML can be used in SELECT, and SELECT INTO statements wherever rowset providers such as table, view or OPENROWSET can appear. To write queries using OpenXML, an internal representation of the XML document must be created using the **sp_xml_preparedocument** stored procedure. This stored procedure returns a handle to the internal representation of the XML document. This handle is then passed to OpenXML, which provides rowset views of the document based on XPaths.

Two XML features in SQL Server 2000 will be made available after the product's release. XML updategrams and XML bulk loading will be included in a Web release shortly after the release date. XML updategrams allow Web developers to use XML to insert, update and delete data from SQL Server 2000 tables. The XML bulk load facility enables bulk loads of data packaged with XML. Both updategrams and bulk load leverage the bi-directional XML-to-relational mapping provided by annotated schemas.

Easy and secure access to data via the Web

SQL Server 7.0 provided features to generate HTML pages from SQL Server data, post SQL Server data to HTTP and FTP locations and send and receive email. SQL Server 2000 goes far beyond this to deliver new and enhanced features for accessing data stored in SQL Server databases and OLAP cubes via the Web. These features ensure that data is available via intuitive user interfaces and fully searchable, even by novice users.

Secure HTTP connectivity to the database

SQL Server 2000 enables URL-based access via several different mechanisms, all of which rely on an included ISAPI DLL that allows users to define virtual roots in Microsoft Internet Information Services (IIS) that are associated with an instance of SQL Server 2000. The URL can be specified to:

- Execute SQL queries directly. URL queries have the form: http://server/vroot?sql="..." For URL queries, Internet applications can compose URL strings that reference a SQL Server 2000 virtual root and that contain a Transact-SQL statement. The Transact-SQL statement is sent to the instance of SQL Server 2000 associated with the virtual root, and the result is returned as a standard rowset. If the FOR XML clause is specified, an XML document is returned instead. Single row, single column queries are supported.
- Access database objects, such as tables, directly. Direct queries have the form: http://server/vroot/dbobject/xpath. For direct queries via HTTP, data is not returned as XML, enabling the direct retrieval of objects like images from the database. The XPath syntax required in this case treats tables (and views) as elements and columns as attributes.
- Execute template files. Template queries have the form: http://server/vroot/vname?params and directly reference a template file, which is a valid XML document consisting of one or more SQL statements. When a template file is specified at the URL, the SQL commands stored in the template file are executed. The queries are replaced with the results and the entire XML document is returned to the requestor.
- Execute XPath queries via an XML View. XPath queries of XML Views have the form: http://server/vroot/vname/xpath?params XPath

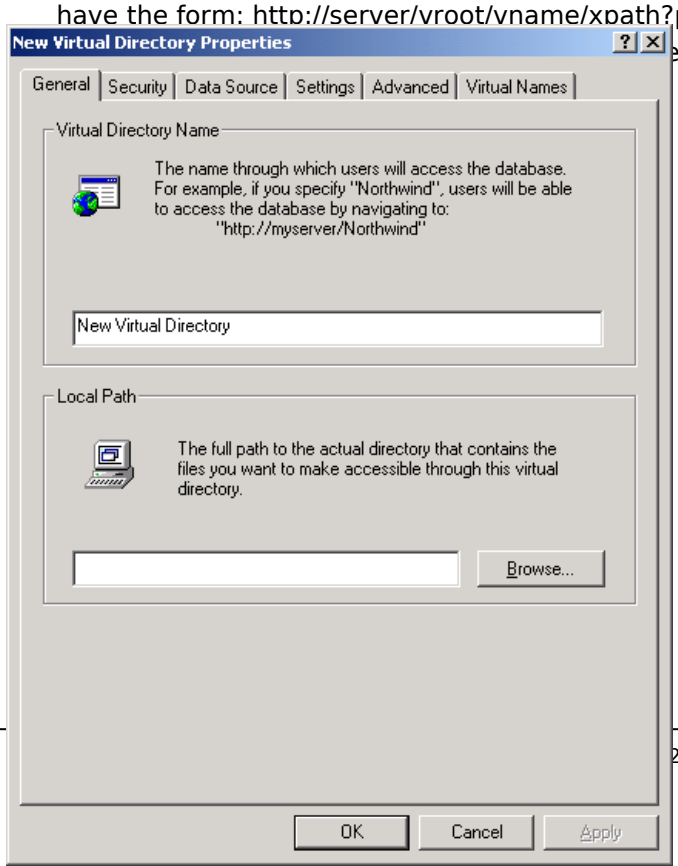
New Virtual Directory Properties
General
Security
Data Source
Settings
Advanced
Virtual Names
Virtual Directory Name
The name through which users will access the database. For example, if you specify "Northwind", users will be able to access the database by navigating to: "http://myserver/Northwind"
New Virtual Directory
Local Path
The full path to the actual directory that contains the files you want to make accessible through this virtual directory.
Browse...
OK
Cancel
Apply

es through XML View database tables.

Figure 1: The virtual root used to access SQL Server 2000 can be specified from the New Virtual Directory Properties dialog.

With all of these options, security is set up per virtual root and permissions are set on SQL Server. The ISAPI DLL itself has three authentication options:

- Anonymous access with Windows or SQL Server ID/password set directly on virtual root and used for all users.
- Basic authentication through clear text transmission of SQL Server id/password (should be used in conjunction with SSL).
- Integrated using Windows ACLs.

Full-Text Search

A very large portion of digitally stored information is in the form of unstructured textual data, both in plain text files and formatted documents. While by far the bulk of this textual data is stored in file systems, many organizations are now storing such data in relational databases. By doing so, the advantages of a database, including high performance access, query capability, simple application-based user interfaces for end users and secure remote access are made broadly available. Unfortunately, most database vendors charge extra for text search features, so use of this functionality is not as widespread as it could be. Other users have had to resort to buying expensive third party offerings to address their needs. These solutions usually involve pulling data out of the database via bridges or gateways and storing them as files in the file system so that textual indexing can be applied. This primitive, brute force approach offers no seamless way for a user to combine a textual query with a regular structured relational query. As introduced in SQL Server 7.0, Full-Text Search allows developers to use and manage structured and unstructured data in a uniform manner.

SQL Server 2000 expands on the integrated Full-Text Search capabilities offered in SQL Server 7.0, allowing users to search through text stored in the database even when the text is contained in a formatted document. Full-Text Search is core functionality in SQL Server 2000 and is part of the default SQL Server install; users no longer need to run a custom setup to

enable Full-Text Search. Moreover, it is included in the Enterprise, Standard, Personal, Developer, and Enterprise Evaluation Editions of the product at no additional charge.

As a brief background, the Full-Text Search capabilities in SQL Server enable:

- Full-text queries on plain text data stored in relational tables
- Full-text queries that are seamlessly integrated into the T-SQL language; single queries can combine full-text and "traditional" search paradigms
- Searching across all indexed columns in a table
- Programmatic access to information about the query (relative rank of hits, weighting by search terms, etc.)

SQL Server 2000 introduces support for Full-Text Search on formatted documents (e.g. Office and HTML documents). For Full-Text Search, documents are stored in BLOB columns in the database and registered for full-text indexing. Document filters enable the extraction of textual content from specific document formats. SQL Server examines each document in order to determine which of several filters is appropriate. At indexing time, the appropriate filter is loaded and passes back the text information necessary to create a full-text index, filtering for noise words.[6] SQL Server 2000 ships with filters for HTML files, text files and Office documents, but developers can write their own filters via the IFilter specification.[7]

Maintenance of full text indexes is critical so that users can find the information they are looking for, even when it has been recently added to the database. In SQL Server 7.0, it was necessary to schedule batch processing jobs to update full-text indexes. SQL Server 2000 introduces change tracking as a new index maintenance technique. Change tracking maintains a list of all changes to full-text indexed data. These changes can be propagated to the index automatically or on a schedule by using the SQL Server Agent.[8] The updating process proceeds asynchronously in the background, providing simpler administration, reduced processing (as a result of more efficient updating), and more up-to-date full-text indexes. Change tracking reduces the need for full and incremental populations in most cases, but when full and incremental population operations are required, SQL Server 2000 performs them more quickly than SQL Server 7.0.[9]

---

[6] To prevent a full-text index from becoming bloated with words that do not aid searching, noise words, such as *a*, *and*, *the*, etc. are ignored. Noise word lists for various languages are provided. These noise word lists can be modified for specific environments.

[7] IFilter is documented at http://msdn.microsoft.com/library/

[8] The mechanism that supports auto-propagation in non-transactional and there is some delay before changes are reflected in the full-text index.

[9] Change tracking is the recommended mechanism, even on near 100%

Other enhancements to the Full-Text Search capabilities of SQL Server 2000 include:

- Support for full-text indexing in failover clusters, allowing full-text indexes to failover in conjunction with the failover of a SQL Server 2000 instance.
- Full-text queries against views for greater flexibility.[10]
- Top_N_By_Rank, to return the top N matches by rank, useful for enhancing performance.
- Multilingual support, as a result of support for column level collations in SQL Server 2000. Word-breaking and verb conjugation are performed intelligently (i.e. they are language sensitive).

Security

Understanding the sensitive and vulnerable nature of Web-based e-commerce applications, SQL Server 2000 introduces significant new security enhancements, not only offering the highest level of security available in the industry, but also making it much easier to achieve that level. To start with, SQL Server 2000 installs with a much higher level of default security, taking advantage of Windows 2000 integrated security out of the box. This makes server lock down in production environments easier and faster.

SQL Server 2000 also introduces a collection of sophisticated new security features: powerful and flexible role-based security for server, database and application profiles, integrated tools for security auditing (tracking 18 different security events and additional sub-events), support for sophisticated file and network encryption (including SSL), as well as Kerberos and delegation. SQL Server 2000 is in the process of being certified under the U.S. Government's C2-level security certification, this process is expected to be complete shortly after SQL Server 2000 is released.

English Query

Traditionally, developers have built applications that hide the generation of queries from the end user either by using algorithms to build them or by offering a static set of queries accessible through an intuitive application user interface. These approaches do not allow for full leverage against data stored in the database, however. An ideal mechanism allows for free-form queries to be generated by novice end users, allowing the developer to focus on other parts of the application rather than building a large number of queries into the application and hoping that they actually meet the needs of end users. Today, with the popularity of the Web, the need for end users to generate free form

modifications of full-text indexes.

[10] In SQL Server 2000, this is supported against views with one underlying base table.

queries is even more critical. Users expect Web sites to offer comprehensive, easy-to-use searching capabilities.

Microsoft English Query enables database developers and administrators to offer all types of end users easy access to data stored in relational databases while supporting power users and those requesting complex data subsets. English Query allows end users to pose questions in English instead of forming a query with an SQL statement. English Query acts as an intermediary, turning the user's question into a proper SQL query that pulls the desired results from the database. English Query is fully integrated into SQL Server 2000 and is included in the Enterprise, Standard, Personal, Developer, and Enterprise Evaluation Editions. Microsoft is unique in the industry in offering such functionality, let alone including it with the core database product.[11]

The creation ("authoring") and deployment of English Query applications has been vastly simplified. Authoring is now hosted in the Microsoft® Visual Studio® version 6.0 development environment, which is included with SQL Server 2000.

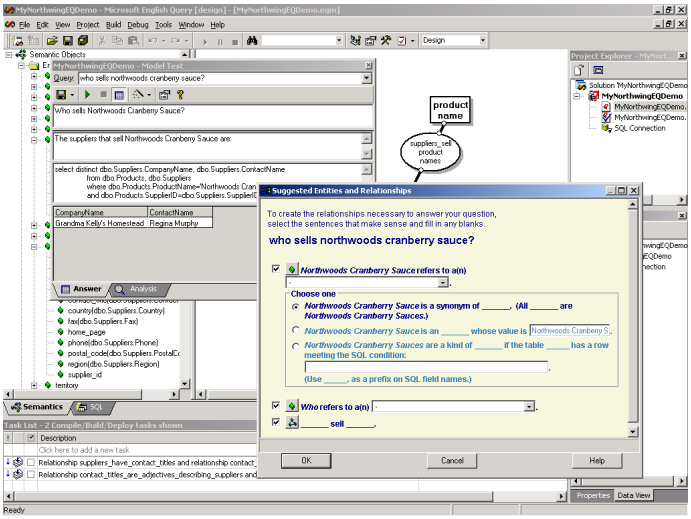

Figure 2: English Query authoring is now performed in the Visual Studio environment.

When beginning an English Query project in this environment, developers can take advantage of the new project wizard to automatically create entities and relationships for all tables in the database. For each entity, the project wizard automatically builds name and trait relationships resulting in automatic creation of approximately 70 percent of the entities and relationships needed in a model. This environment also

[11] Teri Palanca of Giga Information Group: "[Plain-text English querying is] not something we see from other vendors, to have it packaged into the server. [Microsoft is the only company that has] the natural language capability; everyone else is using SQL." Quoted in PCWeek online, "SQL Server Learns English," 3/20/2000.

includes a graphical diagramming tool that displays the entities and relationships in the English Query model being created and that supports drag-and-drop of entities onto one another to create relationships between them automatically. To achieve fine-grained control, developers can use the new, XML-based Semantic Modeling Format (SMF) to persist English Query model information. Used with the Authoring Object Model, SMF provides programmatic authoring of English Query models. Several sample EQ projects are included to help developers quickly learn how to set up their projects.

Deployment of English Query applications to the Web is now as simple as a few clicks. Similarly, testing and maintenance of English Query applications has been simplified via a wizard that can make suggestions for new relationships and entities based on failed user questions. Failed questions are logged and the suggestions English Query offers reflect the changes needed to enable those questions to succeed.

When used in conjunction with Analysis Services (formerly OLAP Services) in SQL Server 2000, developers can now target English Query applications at OLAP cubes via the Multidimensional Expression (MDX) generation capabilities of English Query. A project wizard specifically designed for OLAP is included; it provides an even higher percentage of automatic entity and relationship creation than the aforementioned standard project wizard. Besides integration with OLAP cubes and analysis results, English Query is now able to connect to a variety of data stores via OLE DB, ensuring that integrating data from heterogeneous sources can go hand-in-hand with making it easy to access.[12] Rounding out these integration capabilities, English Query in SQL Server 2000 can generate other forms of SQL in addition to Transact-SQL, including SQL for Oracle and Microsoft Jet SQL.

End users will benefit from integration with Full-Text Search (by generation of SQL queries containing the CONTAINS and FREETEXT predicates), allowing English Query applications to search both structured and unstructured textual data. Sophisticated end users may use the new Question Builder to build queries graphically and retrieve information about an English Query model, including what English phrases can be used to ask about relationships. Question Builder features a three-pane interface that makes the task of determining what information is available in the database much easier, graphically displaying the entities and relationships present in a model.

Powerful, flexible Web-based analysis

SQL Server 2000 delivers several significant features in Analysis Services that allow companies to derive additional value from their data by

[12] Previous implementations of English Query used ODBC only.

enabling fast, flexible analysis. Integrated Data Mining is new to SQL Server 2000 and is a core component for delivering a complete, end-to-end analysis solution in SQL Server. Linked cubes and HTTP access to cubes expand analysis outside the bounds of a company and its intranet to enable new markets for multidimensional data and new ways to browse that data over the Web. Besides these obviously significant features, Analysis Services also includes features, like DISTINCT COUNT, that simplify the life of the analyst. Particularly useful for analyzing user traffic on Web sites, DISCOUNT COUNT has traditionally been a high-end analysis feature because it is relatively difficult to implement. Analysis Services in SQL Server 2000 delivers DISCOUNT COUNT to all users as a new type of measure and allows analysts to answer important questions like “How many unique users hit my site today?” This is just one small example of how the SQL Server team has improved the power and flexibility of Analysis Services in SQL Server 2000.

Data Mining and Business Internet Analytics

Integrated data mining is a new offering for SQL Server 2000 and is delivered as part of Analysis Services in the Enterprise, Standard, Personal, Developer, and Enterprise Evaluation Editions. Data mining technology helps users analyze data in relational databases and multidimensional OLAP cubes to uncover patterns and trends that can be used to make predictions. The data mining capabilities in SQL Server 2000 are tightly integrated with both relational and OLAP data sources. In fact, the results of data mining can be used to create additional cube dimensions for further OLAP data analysis. And you can use them with relational databases by simply issuing SQL queries. The data mining features included in SQL Server 2000 Analysis Services are incorporated in an open and extensible implementation of the new OLE DB for Data Mining specification.

SQL Server 2000 includes two classes of data mining algorithms developed by Microsoft Research: Microsoft Decision Trees and Microsoft Clustering. The Microsoft Decision Trees algorithm actually consists of four different algorithms and is based on the notion of classification. The algorithm builds a tree that will predict the value of columns based upon the other columns in the training set (i.e. a fact table). The decision on where to place each node in the tree is made by the algorithm, and the most significant and differentiating attributes are shown closer to the root of the decision tree. An implementation of the Microsoft Decision Trees algorithm might be used to identify individuals who are most likely to click on a particular banner ad or buy a specific product from an e-commerce site. The Microsoft Clustering algorithm uses a nearest neighbor method to group records into clusters that exhibit some similar, predictable characteristic. Often, these characteristics may be hidden or non-intuitive. For example, the Microsoft Clustering algorithm might be

used to assess customer purchase behavior by age. Of course, Integrated Data Mining in SQL Server 2000 supports algorithms developed by third parties.

Support for data mining is pervasive through Analysis Services and other aspects of SQL Server 2000. New wizards, editors and other user interface elements are provided to simplify the design, creation, training, and browsing of data mining models. Data mining results can he incorporated into OLAP cubes, and MDX syntax has been extended in SQL Server 2000 to assist in the programmatic manipulation of data mining models in connection with OLAP cubes.

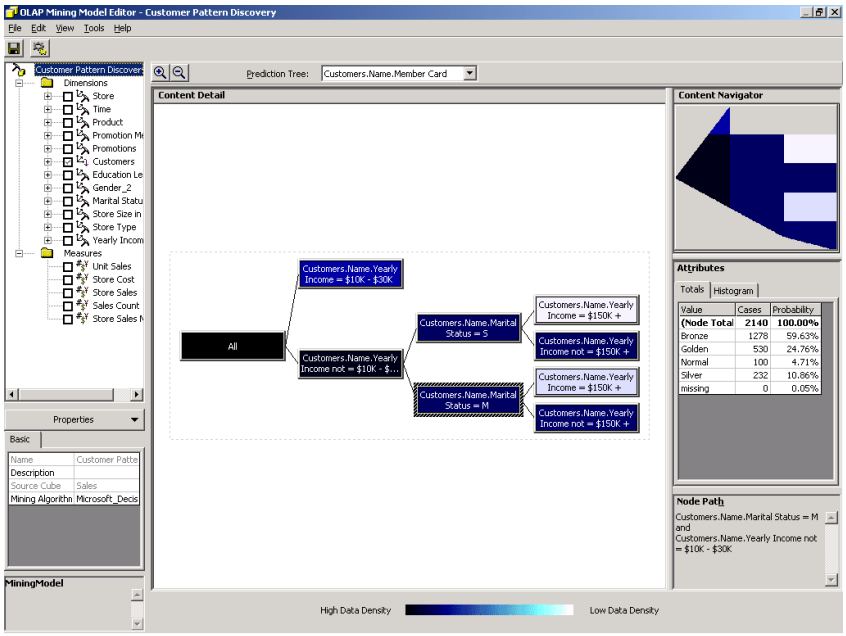

Figure 3: The Data Mining Model Editor displays the results of data mining analyses (here using the Microsoft Decision Trees Algorithm).

Integrated data mining is a key element of Microsoft's strategy to deliver closed- loop Business Internet Analytics in Windows DNA 2000. Closed-loop Business Internet Analytics involves:

1. Collecting information from the online behavior of customers as they browse and search a Web site.
2. Analyzing that information to uncover trends and make predictions (using data mining).
3. Personalizing ads and content for users based on this analysis, such as displaying appropriate banner ads to cross-sell products.
4. Driving decisions back into the operational systems using OLAP Actions.

Ideally, this is a continual process for e-commerce sites that keeps them in tune with their customers and helps them differentiate themselves from their competition. Integration with Commerce Server 2000

simplifies data collection for user clickstreams, transactions, purchase histories and other customer activity data that builds the holistic view of the business activity on a site. This integrated data then enables business managers to modify or create new marketing programs, promotions, and advertising campaigns as well as drive merchandising and site personalization. This information can also be integrated with ERP and CRM systems to enable supply-chain responses to customer behavior.

Linked Cubes and HTTP access to cubes

As companies gather ever-increasing amounts of data on their customers, they seek new ways to analyze and create value from this data. Analysis Services in SQL Server 2000 introduces linked cubes and HTTP access to cubes, two technologies that Web-enable analysis so that users can leverage cubes owned by partners or offered for sale by research firms.

Linked cubes are cubes that are defined and stored on other analysis servers, including servers outside the corporate firewall. To end users, linked cubes appear and function like regular cubes. Linked cubes allow data providers to create, store, and maintain a cube on one analysis server while the cube is also available as one or more linked cubes on multiple analysis servers; data sources using HTTP and HTTPS are supported. Linked cubes use the aggregations of their source cubes and have no data storage requirements of their own. This technique allows one organization to maintain ownership of and update the cube while making the power of the cube available to many consumers simultaneously. It also offers the data provider peace-of-mind; sensitive information can be stored in data sources and cubes on secured servers but still made available broadly through other servers as linked cubes. In addition, by using the Virtual Cube feature to combine an external linked cube from a market research firm with an internal cube containing your sales data, a new level of analytical insight can be gained through immediate evaluations of market share, penetration, relative trends, market coverage and market potential.

Allowing end users to analyze cube data flexibly, in a variety of clients, is also critical to Web-enabling analysis. In SQL Server 2000, Analysis Services uses an HTTP listener built into the server to provide cube access over HTTP. This enables organizations sharing cubes or accessing remote cubes to do so securely over HTTP and through a firewall without opening dedicated ports on a web server.

Secure application hosting

More and more companies are choosing to outsource their applications—from e-commerce to human resources to online company stores—to

Application Server Providers (ASPs). ASPs thrive on economies of scale; in order to provide the best value for their customers they must make best use of their expensive hardware resources. As a result, ASPs often host multiple applications from different companies on a single server. When those applications rely on a database, that means having multiple databases installed. SQL Server 2000 now supports multi-instance installations, providing greater flexibility to ASPs and others with tightly constrained hardware budgets. Multi-instance support allows multiple, independent installations of SQL Server on the same machine. This increases reliability in the case that a misbehaved application degrades the performance of or crashes a database. With the isolation provided by multi-instance, only that application is affected, other applications on top of other instances of SQL Server 2000 keep running. Multi-instance support also enables development and test teams to use a single server for testing different versions of an application under development.

SQL Server 2000 also includes a new method for distributing a disk image of an installation. When an installation is first created, it is marked as a new installation. This functionality allows ASPs and others to install SQL Server 2000, stop the server, clone the disk image, and then distribute it as needed. On the first startup of the distributed server it can be renamed. This functionality simplifies the creation of standard or default databases in server farms.

HIGHLY SCALABLE AND RELIABLE

SQL Server 7.0 introduced the most modern database architecture in the industry. SQL Server 2000 builds on that modern architecture by offering scalability and reliability for even the most demanding enterprise applications (on affordable and manageable hardware from a variety of vendors). As proof of this, SQL Server 2000 has achieved world-record performance for a number of industry benchmarks—eclipsing competing products regardless of the hardware and operating system those competitors were running. In addition, limitations on performance and scalability as a result of hardware constraints are quickly vanishing. This ensures a clear roadmap to support future growth for those customers building on the Windows DNA platform with SQL Server.

One of the notable benefits of Windows DNA is its embrace of the concept of software scale out. Windows DNA has already delivered software scale out technologies at the web and applications tiers, allowing application code to seamlessly exploit inexpensive, commodity hardware added incrementally to server farms. Until now, software scale out has not been enabled “out of the box” by Microsoft solutions on the data tier. SQL Server 2000 delivers the first major installment of the Microsoft vision for software scale out on the data tier through a feature called Distributed Partitioned Views (scale up techniques are also fully supported by SQL Server).

Longer term, there are additional significant technology milestones ensuring an accelerating expansion of scalability for customers building and deploying solutions on the Microsoft platform. These innovations include the expected introduction in mid-2000 of Intel Pentium III Xeon-based servers with 16 and 32 processors, offering many times more power than existing eight-processor systems. In addition, following the introduction of Intel’s 64-bit Itanium platform later in 2000, Microsoft will release 64-bit versions of Windows 2000 and SQL Server 2000. These state-of-the-art systems will allow customers to utilize terabytes of system memory for the most complex applications.

Scalability and reliability for E-Commerce solutions

SQL Server 2000 is capable of supporting the growth requirements of any Web site, as well as the largest enterprise systems. It includes Distributed Partitioned Views and log shipping to distribute workloads and takes full advantage of SMP hardware with new parallel operations.

Distributed Partitioned Views

Distributed Partitioned Views shares the database processing load across a group of servers by horizontally partitioning the SQL Server data. These servers cooperate in managing the partitioned data but operate

autonomously. The partitioning of data is transparent to applications accessing the database; the application "sees" a full copy of all tables no matter what server in the federation it accesses. All servers accept connections and process both queries and updates, distributing scans and updates as needed. The SQL Server 2000 query processor contains a number of enhancements that make these views updateable, and lets them perform well if they need to use distributed queries to pull data from other member servers.

Distributed Partitioned Views is made possible through the implementation of shared-nothing clustering technology. The shared-nothing approach to clustering is generally considered to be superior to the shared-disk approaches because single points of failure and bottlenecks as a result of shared resources do not occur. In SQL Server 2000, shared-nothing clustering is implemented as a "federation" of independently administered servers that cooperate to jointly manage workload; this is a tremendous step toward a centrally managed shared-nothing offering, and affords the same level of performance.

More information on Distributed Partitioned Views is available in the "Microsoft SQL Server MegaServers: Achieving Software Scale Out" white paper at: msdn.microsoft.com/sqlserver

Log Shipping

Log Shipping is fully integrated into SQL Server for the first time in SQL Server 2000. Log Shipping can be used to provide "warm" standby on multiple back up servers by automatically feeding transaction logs from one database to another on a continual basis. Continually backing up the transaction logs from a source database and then copying and restoring the logs to destination databases keeps the destination databases in synchronization with the source database. This can improve scalability by offloading query processing from the main machine (the source server) to read-only destination servers. Moreover, it aids reliability by providing warm standby servers.

SMP-related features

SQL Server 2000 does more operations in parallel to take greater advantage of increasingly common Symmetric Multiprocessing (SMP) hardware, from dual processor systems prevalent even in small businesses to 16- and 32-way systems in the data centers of Fortune 500 companies. Parallel index creation is enabled by building sub-indexes for specific ranges of an index. Separate threads (running on separate processors if available) build these sub-indexes, fed by parallel scans. When the sub-indexes are completed they are combined into a complete index by a coordinating thread. This parallelism is particularly valuable for highly transactional databases that, because they are frequently updated, require routine index rebuilds. Similarly, multi-terabyte data

warehouses benefit from this feature when creating an index on a fact table that might otherwise take hours to complete.

The DBCC utility has been enabled for multi-processor systems as well. It runs with parallel threads, so the speedup is dependent on the number of CPUs present. Other enhancements to DBCC further improve performance; higher concurrency is enabled because SQL Server 2000 takes a Schema Stability Lock, rather than a shared lock, on tables it is checking. For more information on parallel operations in SQL Server 2000, see Books Online.

Scalability and reliability for Line-of-Business solutions

When a company has invested in an application that is truly mission critical, it must be assured that the application will scale to meet the company's needs and stay up and running. SQL Server 2000 includes enhanced failover cluster support to keep applications up and running even when disaster strikes. It can also address 64 GB of RAM and take advantage of 32 CPUs. To speed data between servers, SQL Server 2000 introduces built-in support for SANs (system area networks) for high-speed cluster interconnects.

Enhanced failover clustering

Failover clusters are both easier to manage and more flexible in SQL Server 2000 Enterprise Edition. The installation, creation, configuration and maintenance of failover clusters are now more accessible from setup and SQL Server Enterprise Manager. SQL Server 2000 setup automatically detects the presence of Microsoft Cluster Services, so no separate cluster wizard is needed. Nodes can even be added or removed during setup. From Enterprise Manager, reinstallations and rebuilds can proceed on any node without affecting the others in the cluster. Moreover, failover clustering now works with replication, Distributed Partitioned Views and other applications that are dependent on server names.

Operational enhancements allow failover and failback to or from another node in a cluster. Using Windows 2000 Advanced Server in an "Active-Passive" configuration, an instance of SQL Server 2000 is run on a primary machine while a secondary instance on a second machine is idle until failover. In an "Active-Active" configuration, SQL Server 2000 runs multiple servers simultaneously with different databases, allowing for organizations with more constrained hardware requirements (i.e. no designated secondary systems) to enable failover to or from any node without having to set aside hardware. In addition, with Windows 2000 Datacenter Server, SQL Server 2000 Enterprise Edition supports four-node failover clusters. When one node fails, the SQL Server and OS resources can failover to any other surviving node.

Large memory and SMP support

On Windows 2000 Datacenter Server, SQL Server 2000 extends its reach to 64 GB of RAM and up to 32 CPUs. This represents thorough support for scale-up scenarios and can be used in conjunction with scale out techniques to handle the largest data sets and transactional loads.

VI SAN support

The VI SAN (Virtual Interface System Area Network) support in SQL Server 2000 allows SQL Server to communicate directly with devices connected via a SAN to transfer high volumes of data or transactions with low latency. Microsoft has worked closely with Giganet (cLan) and Compaq (Servernet 2) to offer this direct access to SAN devices.

Scaling applications to new classes of devices

SQL Server 7.0 provided notebook- and laptop-based solutions with seamless scalability up to SQL Server 7.0 Enterprise Edition. SQL Server 2000 Windows CE Edition further extends the SQL Server product family to new classes of devices targeting mobile and embedded solutions for sales force automation, service dispatch, point-of-sale, shop floor management, and even set-top boxes.

SQL Server 2000 Windows CE Edition provides developers with a consistent programming model and API set. Developers can quickly and easily develop applications for SQL Server 2000 Windows CE Edition by taking advantage of their existing experience with the Visual Basic® and Visual C++® development systems and by using a consistent programming model based on OLE DB and ADO.

SQL Server 2000 Windows CE Edition provides all of the essential relational database functionality and the small footprint required for mobile and embedded scenarios. SQL Server 2000 Windows CE Edition supports all of the major CPUs for Windows CE and offers three key components: a robust data store for data persistence; a query processor to simplify and optimize complex queries; and reliable and scalable bi-directional merge replication capabilities to synchronize data between appliances and a centralized SQL Server 2000 database.

Scalability and reliability for Data Warehousing solutions

Data Warehousing and analysis offers its own set of challenges to scalability and performance. Storage and access requirements for multidimensional data grow quickly and facilities for handling complex analyses and very large data sets are critical. SQL Server 2000 includes new features to enable analysis in these scenarios.

Indexed views

Complex reporting scenarios, encountered frequently in data warehousing applications, extract high demands on the database server during querying. When these queries access views, performance can decline as the database dynamically merges the logic needed to build the view result set into the logic needed to build the complete query result set from base table data. The overhead of this operation can be substantial, particularly for views that involve complex processing of large numbers of rows (e.g. large data aggregations or multi-table joins). Because the result set is not stored permanently in the database (with a standard view), future accesses to the view are likely to incur the cost of building the result set each time the query is executed.

SQL Server 2000 enables the performance of queries accessing such

views to be dramatically improved by allowing the creation of a unique clustered index on the view. When such an index is created, the view is executed and the result set is stored in the database in the same way a table with a clustered index is stored. This effectively *materializes* the results of the query in the database. The SQL Server query optimizer will use the view index in queries that do not directly name the view in the FROM clause. Existing queries can benefit from the improved efficiency of retrieving data from the indexed view without having to be recoded. For certain types of views, this can improve performance exponentially.

From an ease of use perspective, the implementation of indexed views in SQL Server 2000 is unique in the industry. Competing products that implement indexed or materialized views require queries to be specifically coded to reference the views, so existing queries must be updated with hints to force the query processor to recognize indexed or materialized views. As mentioned above, the implementation of indexed views in SQL Server 2000 does not require any recoding of queries – the SQL Server query processor automatically recognizes and leverages indexed views when appropriate without any DBA intervention.

Distributed Partitioned Cubes

SQL Server 2000 extends software scale out to data warehousing solutions through Distributed Partitioned Cubes. In order to attain flexible data storage and additional query performance, developers can use the Partition Wizard to easily scale out one logical cube into separate physical partitions on multiple servers with complete transparency.

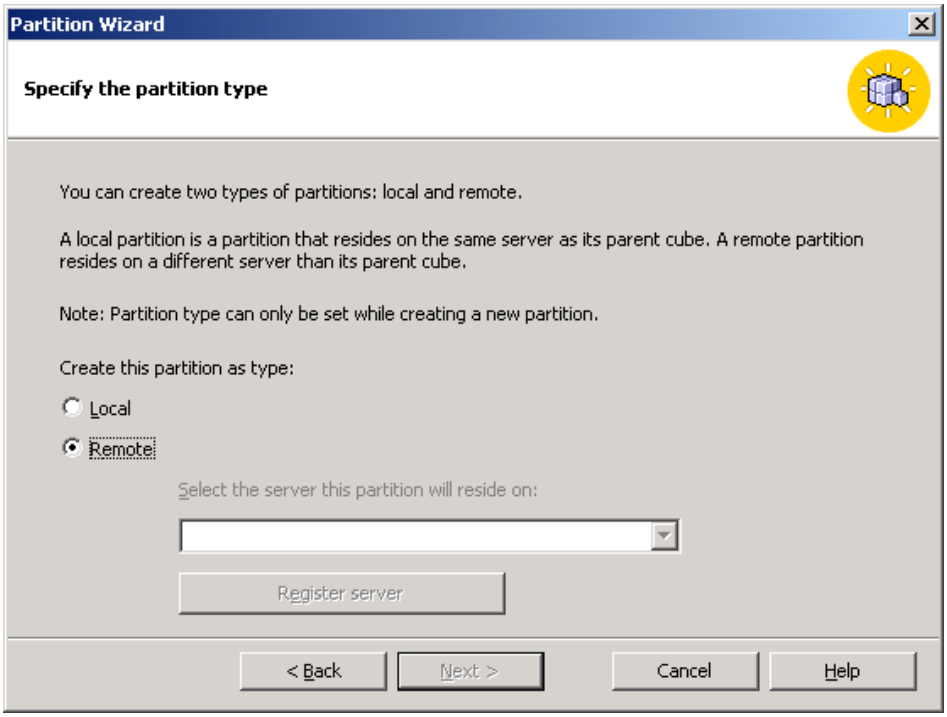
Partition Wizard
Specify the partition type
You can create two types of partitions: local and remote.
A local partition is a partition that resides on the same server as its parent cube. A remote partition resides on a different server than its parent cube.
Note: Partition type can only be set while creating a new partition.
Create this partition as type:
Local
Remote
Select the server this partition will reside on:
Register server
< Back
Next >
Cancel
Help

Figure 4: Setting up Distributed Partitioned Cubes is simple with the new Partition Wizard.

Large data set analysis

Analysis Services in SQL Server 2000 supports MOLAP (Multidimensional OLAP) dimensions in the range of tens of millions of members by utilizing up to 64 GB of memory. Through the addition of ROLAP (Relational OLAP) dimensions, which are dimensions stored as relational tables, Analysis Services supports dimensions on the order of hundreds of millions of members. ROLAP storage mode causes aggregations to be stored in relational tables that are not subject to the size limitations of MOLAP. In some cases, indexed views are created instead of tables if the partition's source data is stored in SQL Server 2000 and if certain criteria are met. The Partition Wizard can be used from Analysis Manager to easily separate one logical cube into separate physical partitions.

In contrast to the MOLAP storage mode, ROLAP does not cause a copy of the source data to be stored; the partition's fact table is accessed to answer queries when the results cannot be derived from the aggregations or client cache. With the ROLAP storage mode, query response is generally slower than that available with the other two storage modes, MOLAP and HOLAP (Hybrid OLAP). Typical applications of ROLAP include large data sets that are infrequently queried (such as historical data) or data sets that are too large to be stored in MOLAP or HOLAP mode.

Maximized uptime and reliability

While scalability is important, no amount of power or performance is useful if a database isn’t up and running. Therefore, maximizing uptime is essential. In SQL Server 2000, more operations can be performed online so that the server can stay up and running. Additionally, new backup features ensure that backups can be performed with little or no impact on server performance or availability.

More online operations

In databases that are highly transactional or are otherwise updated routinely, the uptime implications of index reorganizations have often forced database administrators to perform these operations during non-peak times. With SQL Server 2000, index reorganization can be performed online with little impact on performance, even for demanding OLTP applications. Furthermore, online index reorganization uses minimal data space, is fully logged (for recovery) and may be stopped and restarted.

Differential database backups

SQL Server 2000 includes the capability to perform differential backups.

A differential backup makes a copy of all the pages in a database that have been modified since the last database backup. Using this approach, backups run relatively quickly and are smaller in size than other types of backups. Moreover, differential backups may be performed while users access the database. Because of their speed and low server impact, differential backups can be made more frequently than other types of database backups, decreasing the risk of data loss.

Server-less snapshot backups

Another new feature in SQL Server 2000 that results in higher percentage uptime while reducing impact on resources is server-less snapshot backup. Server-less snapshot backups are functionally equivalent to full database or file/filegroup backups, but can be performed with no impact on the performance of the server (hence the use of the term “server-less”).

Server-less snapshot backups require the use of a third party virtual device interface (VDI) application that can communicate directly with an advanced enterprise storage system supporting split mirror or copy-on-write operations. A common scenario for using server-less snapshot backups would be in a situation that included 3-way mirroring. The VDI backup application can “break off” one of the mirrors (two mirrors continue to support users) and either backup that mirror to tape or make it available to another system. In the latter case, the backup effectively becomes an immediately available standby database. Server-less snapshot backups can be restored very quickly when brought online from a (disk-based) mirror. Not surprisingly, restores from tape take longer.

Server-less snapshot backups offer great flexibility to organizations. Besides being a technique for initializing “warm” standby servers, they can be used to easily create reporting or test databases, with effectively no impact on production servers. SQL Server maintains the history of server-less snapshot backups in MSDB, and they can be rolled forward using conventional differential and log backups.

FASTEST TIME-TO-MARKET

SQL Server is the database of choice when time-to-market is a critical factor. Ease of management, excellent programmability, deep tools integration and integrated services have helped SQL Server become the preferred database for the Web, where delivering a solution before the competition is necessary for success. SQL Server 2000 continues to build on its repertoire of achievement in this area by delivering enhanced management tools that automate management and tuning, tools and utilities that speed development, debugging and data transformation, an expanded set of integrated services for analysis, and thorough integration with the other servers making up the Windows DNA 2000 platform.

The SQL Server strategy for enabling management and application development is to provide tools that are "simple but not simplistic." This means providing a broad range of consistently designed, highly usable tools and interfaces. Fine-grained control is there if the developer or database administrator wants it, but the product is designed to be scalable, reliable and self-managing out-of-the-box.

Simplified management and tuning

SQL Server 2000 delivers technology that brings database applications to market more quickly and ensures that the cost of managing and tuning them is minimized throughout their lifetime. As database administrators (DBAs) gain responsibility for more and more databases each year, centralized management becomes increasingly important. SQL Server 2000 enables centralized management but also endeavors to automate management and tuning as much as possible to reduce the burden on the administrator. When automation is not an option, SQL Server 2000 offers intuitive wizards to quickly step administrators through complex tasks.

Integration with the Active Directory™ service

The integration between SQL Server 2000 and the Active Directory service in Windows 2000 allows SQL Server databases to be managed centrally alongside other enterprise resources, vastly simplifying system management in large organizations.

Servers and their attributes are registered automatically in the Active Directory on SQL Server startup, and administrators can search and locate servers through the Active Directory Search feature in SQL Server 2000. For example, a user might use the directory to locate all of the servers running one or more instances of SQL Server 2000 with a particular database name installed. The search capabilities also allow administrators to easily answer questions like "What SQL Servers are on

the network?" Besides enumerating databases, Active Directory also lists replication publications, cubes and data mining models (entries are made each time one of these items is created). These entries can then be browsed using Explorer or a custom application. In addition, Active Directory provides database location transparency – aliases remain the same even when servers are moved to different machines or locations.

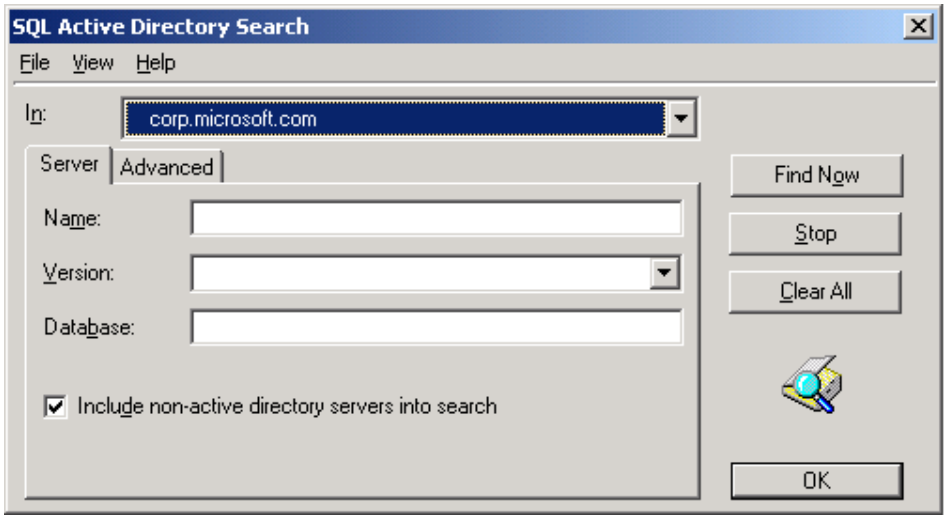

Figure 5: Users can search the Active Directory for SQL Servers.

Dynamic self-management and tuning

Databases must be tuned in order to achieve optimal end-user performance, but there is no single 'correct' way to tune a database that satisfies all end-users, so databases are generally tuned to match the specific usage patterns of the data. The biggest challenge facing DBAs in this area is that the usage patterns of the data change frequently as new tools and new users access the data in different ways. This can make it difficult for DBAs to maintain a database in a 'tuned' state for long periods of time without continual attention.

SQL Server performs several operations to keep itself tuned; in SQL Server 2000, several formerly static configuration parameters have been replaced with dynamic algorithms that rely on adaptive feedback. In particular, several new algorithms have been added to dynamically optimize for recent hardware innovations including large memory, ultra-high-speed disk subsystems and SANs. Additionally, SQL Profiler includes new events for Data File Auto Grow, Data File Auto Shrink, Log File Auto Grow and Log File Auto Shrink. These events allow DBAs to monitor the circumstances under which automated management tasks are occurring. Developers, in turn, can troubleshoot problems by capturing these events on a production system, and replaying them on a test system. Of course, developers can also create applications that use SQL Profiler stored

procedures to automate the adjustment of the operational characteristics of SQL Server.

Copy Database Wizard

The new Copy Database Wizard in SQL Server 2000 enables the simple task of copying and moving databases with minimal server downtime. This feature can be used to set up test environments, move databases between servers or instances, perform hardware migrations and even upgrade from SQL Server 7.0 to SQL Server 2000.

The wizard can be run on source machines, destination machines or non-participating third systems. Regardless of where the wizard is initiated, database copying is implemented as DTS tasks whose packages are executed on the destination machine(s), allowing for flexible scheduling. The wizard supports the copying of global namespace objects, error messages, logins and jobs. Administrators select the databases to be copied, leaving others available to the original (source) server.

Quick development, debugging and data transformation

Delivering robust, scalable database applications quickly requires tools that guide and automate development efforts while also offering powerful, fine-grained control as necessary. SQL Server 2000 offers developer tools to assess and enhance performance, automate the manipulation of data between the widest array of sources and quickly build new applications by taking advantage of existing code.

Query Analyzer

SQL Query Analyzer comes equipped with a T-SQL debugger that allows you to control and monitor the execution of stored procedures. The debugger supports traditional functionality such as setting breakpoints, defining watch expressions, and single stepping through procedures. Developers can use it to view global and local variables and the call stack. It also now adds the capability for server-side statement tracing and client-side statistics to help developers tune and optimize their applications. Through the addition of templates and an object browser, Query Analyzer becomes a true SQL authoring environment for the full range of users from those inexperienced with SQL Server to extremely knowledgeable DBAs.

Data Transformation Services

Data Transformation Services (DTS) has been updated for SQL Server 2000, expanding its abilities to move and transform data from any source. The SQL Server team has focused on improving the usability of DTS while expanding its programmability.

In the realm of programmability, DTS users can now take advantage of a

multi-phase data pump. This means that developers can access and manipulate the operations of the data pump through an expanded number of interfaces and at multiple points during the progression of data transformations. The multi-phase data pump allows for much more flexible handling of transformation and insert errors/failures. When combined with the addition of parameterized source select statements and improved error logging, DTS now allows developers to build restart ability into their DTS packages. Selecting an option in the DTS Designer of SQL Server Enterprise Manager enables the multi-phase data pump UI.

DTS packages can now be saved as Visual Basic® code. This allows an easy learning path to developing DTS packages through programmatic interfaces, and when used in conjunction with Microsoft Visual SourceSafe™ version control system provides an alternative method of version control and backup for packages. Because packages can now call each other when they are executed, the developer has more flexibility in reusing existing packages. New tasks and custom transformation also expand the data manipulation capabilities, performance, flexibility and integration of DTS to products from other vendors.

User-defined functions and programmability enhancements

In addition to the T-SQL functions built-in to the product, SQL Server 2000 allows developers to create user-defined functions. User-defined functions are T-SQL functions that begin with using the CREATE FUNCTION statement. They are subroutines made up of one or more T-SQL statements and can be used to encapsulate code for reuse. Developers can save time by incorporating routinely used logic into user-defined functions as opposed to using stored procedures. The mechanics of calling a user-defined function are identical to those used to call a built-in T-SQL function. User defined functions can return scalar-valued results and, unlike competitive products, table-valued results.

SQL Server 2000 includes other notable programmability enhancements including:

- **Cascading Referential Integrity Constraints** – The actions taken when an attempt is made to delete or update a key to which existing foreign keys point are controllable using the new ON DELETE and ON UPDATE clauses in the REFERENCES clauses of CREATE TABLE and ALTER TABLE statements.
- **INSTEAD OF and AFTER Triggers** – INSTEAD OF triggers are executed instead of the triggering action (for example, INSERT, UPDATE, DELETE). When defined on views, they greatly extend the types of updates a view can support. AFTER triggers fire after the triggering action, first and last firing AFTER triggers can be specified.
- **Indexes on Computer Columns** – Indexes can be defined on computed columns, so long as the expression defined for the column

meets certain restrictions, such as only referencing columns from the table containing the computed column, and being deterministic.

- **New Data Types** – *bigint* is an 8-byte integer type. *sql_variant* is a type that allows the storage of data values of different data types. The new *table* type allows applications to temporarily store results for later use. It is supported for variables and as the return type for user-defined functions.
- **Column Level Collations** – Store objects that have different collations in the same database. Collations can be specified at the database level or at the column level.

Integrated and extensible Analysis Services

The goal for Analysis Services (formerly OLAP Services) in SQL Server 2000 is to provide a complete, end-to-end platform for analysis including relational storage, data extraction, OLAP optimization and querying, data mining, and semantic modeling, among others. Users can mix-and-match algorithms and tools from Microsoft and third parties to customize their analysis applications. Moreover, Analysis Services is designed for database administrators and application developers rather than an elite audience of statistical experts. Anyone with a solid foundation in Structured Query Language (SQL) and Visual Basic should be able to understand, use and programmatically extend the offerings in Analysis Services. Finally, Analysis Services is considered a core offering and is included (at no additional charge) with the Enterprise, Standard, Developer and Enterprise Evaluation Editions of SQL Server 2000.

Since the introduction of OLAP Services with SQL Server 7.0, Microsoft has become a leader in providing the foundation for Business Intelligence. In SQL Server 2000, Microsoft extends its leadership in this space by providing a host of new functionality in Analysis Services. Beyond the core Data Warehousing and OLAP features that were offered in SQL Server 7.0, Analysis Services in SQL Server 2000 provides integrated Data Mining (see Data Mining and Business Internet Analytics above), new graphical tools for building and managing analysis data as well as security and dimension enhancements that improve flexibility.

With Analysis Manager in SQL Server 2000, multiple users can easily and efficiently administer an analysis server. Additionally, remote partitions can be used to distribute a cube's data among multiple analysis servers and administer the cube on a central analysis server (i.e. the server that contains the meta data). Analysis Manager also includes two new graphical components that assist with cube construction. The Virtual Cube Editor allows for virtual cube editing and is similar in function to the Cube Editor. Analysis Services introduces several features to offer flexible security and support for new dimension types. Security is now definable both for dimensions and cells. Roles can be used to control end user access to dimensions, limiting access to individual dimensions, levels, and members, and enabling a variety of read and read/write permissions. Cell security may also be defined by setting role options in Analysis Manager. Roles can be restricted to any combination of a cube's cells.

The following new dimension types are introduced in SQL Server 2000 Analysis Services:

- Parent-child – This dimension type supports hierarchies based on parent-child links between members in columns in a source table.

Examples of such hierarchies include organization charts and part assemblies/bills-of-materials. These dimensions are sometimes referred to as “unbalanced”.

- Ragged – This dimension type has at least one member whose logical parent is not in the level immediately above the member. An example of a ragged dimension would be a country that does not have a state or province level between the country and city levels.
- Changing – This dimension type permits a wider array of changes than other dimensions without requiring a cube to be fully processed after changes. This increases the availability of cubes to client applications.
- Write-enabled – This dimension type can be updated through Analysis Manager and any client applications that supports dimension writeback. Roles are used to control dimension write access from client applications
- Virtual – This dimension type has members that are determined by the members of another dimension.

OLAP Actions

When an end-user has performed an analysis, there is often the need to launch associated applications or link to information. OLAP Actions, a new feature of Analysis Services in SQL Server 2000, allow end users to act upon the outcomes of their analyses to automatically drive business processes. An Action is an end user-initiated operation upon a selected cube or portion of a cube. The operation can launch an application with the selected item as a parameter or retrieve information about the selected item. By implementing Actions, which are easily defined with the Action Wizard, developers can transform client applications from sophisticated data rendering tools to integral parts of a feedback loop in an enterprise operational system. By using Actions, end users can go beyond traditional analysis and initiate solutions to discovered problems and deficiencies.

For a simple example of how Actions might be used, consider an end user at a shoe retailer browsing an inventory cube and noticing that the current stock of white shoes is low. The end user selects the Order Action on the white shoes member. This Action initiates a new order for more white shoes via the order entry system.

Custom rollups

By default, rollups in Analysis Services are additive. Developers and database administrators building analysis solutions need greater flexibility, however. New custom rollup operators provide a simple way to control how member values are rolled up to their parent’s values. Members are tagged with one of the following valid operators: +, -, *, /, and ~. Each of these performs its indicated mathematical action. The

custom rollup operator is applied to the member when evaluating the value of the member's parents. For multiplication and division, the value of the member is multiplied or divided by the aggregate value of the preceding sibling members. Precedence is by member order and the operators themselves are stored in a separate column in the source database. The tilde (~) signifies that a particular member should not be rolled up.

Custom rollup operators are assigned to the name of a column, either when creating them as an optional feature of the new parent-child dimensions in Dimension Wizard or when adding them to existing dimensions via the Dimension Editor or Cube Editor. Custom rollup operators provide functionality similar to that offered by custom rollup formulas, albeit simplified. A custom rollup operator is unique for each level member, while a custom rollup formula applies to all level members.

CONCLUSION

With the most modern architecture in the database industry, SQL Server 7.0 was a breakthrough release for Microsoft. Not only did SQL Server 7.0 offer enterprise-class scalability, reliability and functionality, it also provided a rock solid foundation that has allowed Microsoft to rapidly and continually innovate. In the time since its release, SQL Server 7.0 has become the preferred database of many Customer Relationship Management (CRM), Business Intelligence (BI), Enterprise Resource Planning (ERP), and other line-of-business application vendors and customers because of its performance, scalability, manageability, programmability and value. Additionally, it has achieved great success as an Internet database.[13]

As the next major release of SQL Server, SQL Server 2000 builds upon the modern, extensible architecture introduced in SQL Server 7.0. SQL Server 2000 carries forward this tradition of excellence in database software design and engineering to become the foundation for ongoing innovation in the SQL Server product line, from 64-bit and Windows®-based CE releases expected this calendar year, to "Yukon," the next major release of SQL Server following SQL Server 2000.

Microsoft SQL Server 2000 is the complete database and analysis solution for rapidly delivering the next generation of scalable Web applications. As a core component of Windows DNA, it dramatically reduces the time required to bring e-commerce, line-of-business, and data warehousing applications to market while offering the scalability needed for the most demanding environments.

[13] "Where is Web Data Stored?" Zona Research Quiniela, January 20, 2000. See: http://www.zonaresearch.com/deliverables/quinielas/indexes/2000/jan00.htm

SQL SERVER 2000 NEW FEATURES OVERVIEW[1]

E-Commerce and Internet Features				
Feature	**Benefit and Description**	Enterprise Edition	Standard Edition	Personal Edition
Internet standards support	Integrate existing systems without programming through comprehensive support for W3C standards including XML, Xpath, XSL and HTTP.	☑	☑	☑
XML Views	View and access relational data using XML techniques by easily mapping XML elements and attributes to relational schema.	☑	☑	☑
URL and HTTP access	Flexibly access data from the Web via a URL. Use Structured Query Language (SQL), XML templates or XPath in URLs for querying.	☑	☑	☑
SELECT … FOR XML	Return XML from SQL queries and easily control the shape of the returned XML tree with three formatting options.	☑	☑	☑
OpenXML	Access, manipulate and update XML documents as if they were tables using T-SQL and stored procedures.	☑	☑	☑
XML Updategrams[2]	Use XML to selectively insert, update and delete table data from anywhere, even through firewalls.	☑	☑	☑
XML bulk load[2]	Seamlessly transmit, transform and load data from any source into SQL Server 2000 relational tables.	☑	☑	☑
Full-Text Search	Enable Full-Text Search across the Web and intranets for formatted documents (e.g. Word, Excel, and HTML). Track changes automatically.	☑	☑	☑
English Query	Allow all users to access data via natural language queries. Graphically author queries with wizards in the included Visual Studio environment. Generate MDX to query cubes.	☑	☑	☑
Integrated Data Mining	Analyze collected relational and OLAP data including clickstreams and purchase histories to recognize trends and make predictions. Close the loop by delivering personalized content.	☑	☑	☑
Linked cubes	Link cubes on different servers to supercharge analysis. Securely share or sell source cube data over the Internet.	☑		
HTTP access to cubes	Make cubes available through firewalls for Web-based analysis applications using an HTTP listener built into SQL Server 2000 Enterprise Edition.	☑		
Multi-instance support	Run reliably in hosted scenarios with separate database instances per customer or application.	☑	☑	☑
Security	Protect data with higher default security on install. Includes support for	☑	☑	☑

[2] XML Updategrams and XML bulk load will be available from the Microsoft SQL Server home page (http://www.microsoft.com/sql) as a web release after the release of SQL Server 2000.
[2]

	Secure Sockets Layer (SSL) connections and Kerberos. C2-level certification underway.			
Installation disk imaging	Create standard or default databases for server farms easily and rapidly from any machine on the network with built-in cloning technology.	☑	☑	☑
Scalability and Reliability Features				
Feature	**Benefit and Description**	Enterprise Edition	Standard Edition	Personal Edition
Distributed Partitioned Views	Achieve software scale out on the data tier by partitioning workload across servers. Add additional servers for greater scalability.	☑		
Log Shipping	Automatically keep databases in synch for warm standby on multiple backup servers to share load—no matter how physically far apart.	☑		
Parallel index creation	Take full advantage SMP hardware to speed up index creation, easing the load on frequently updated systems.	☑		
Parallel scan	Accelerate scans on multi-processor systems to improve overall database performance.	☑		
Parallel DBCC	Verify database data quickly and efficiently with multi-processor support.	☑		
Failover clustering	Install failover-ready databases directly from Setup. Use active-passive failover with standby hardware or active-active failover in hardware constrained environments. Databases can failover to any surviving node with 4-node failover.	☑		
Failover cluster management	Reinstall or rebuild any node in a failover cluster without impacting other nodes. Easily configure failover for replication and Distributed Partitioned Views.	☑		
32 CPU SMP system support[3]	Scale-up SQL Server 2000 databases to SMP systems with as many as 32 processors.	☑		
64 GB RAM support[4]	Handle the largest data sets and transactional loads with up to 64 GB of RAM for SQL Server 2000.	☑		
VI SAN support	Speed communications between servers with SQL Server 2000 support for direct communication with System Area Network (SAN) devices.	☑		
Indexed views	Create indexes on views to improve performance of existing queries without recoding. Speed up analysis and reporting that rely on complex views.	☑		
ROLAP dimension storage	Perform analyses on huge dimensions (10M+ members) by storing such dimensions relationally.	☑		

[3] Requires Windows 2000 Datacenter Server.
[4] Requires Windows 2000 Datacenter Server.

Distributed Partitioned Cubes	Easily use the Partition Wizard to separate one logical cube into separate physical partitions with complete transparency for flexible data storage and improved query performance.	☑		
Online index reorganization	Keep the server up and running while reorganizing indexes to improve performance.	☑	☑	☑
Differential backup	Run backups quickly and with low server impact by only backing up changed pages.	☑	☑	☑
Server-less snapshot backup[5]	Backup data seamlessly with effectively no impact on the database server. Restore quickly or create warm standby servers.	☑	☑	☑

[5] Requires third party virtual device interface (VDI) application and enterprise storage system supporting split mirror or copy-on-write operations.

Management and Development Features				
Feature	**Benefit and Description**	Enterprise Edition	Standard Edition	Personal Edition
Active Directory integration[6]	Manage databases centrally alongside other enterprise resources. View and search for servers, replication publications, cubes and more.	☑	☑	☑
Self-management and tuning	Rest easy as SQL Server 2000 uses dynamic algorithms to automatically adapt to constantly changing runtime environments.	☑	☑	☑
Copy Database Wizard	Move and copy databases and objects between servers easily. Schedule migrations through integration with DTS.	☑	☑	☑
SQL Query Analyzer	Debug stored procedures. Set breakpoints, define watches, view variables and step through code. Trace executing code on server or client. Easily write T-SQL based on templates.	☑	☑	☑
Data Transformation Services (DTS)	Program the multi-phase data pump and save DTS packages as Visual Basic code.	☑	☑	☑
User-defined functions	Achieve code reuse by creating T-SQL functions. Incorporate routinely used logic to simplify development.	☑	☑	☑
Cascading referential integrity constraints	Control how changes propagate through tables when keys are updated.	☑	☑	☑
INSTEAD OF and AFTER triggers	Execute code flexibly by specifying what happens in lieu of an operation or after it.	☑	☑	☑
Indexes on computed columns	Define indexes on column types even when the data in the column is computed from other columns.	☑	☑	☑
New data types	Store and reference data flexibly with *bigint*, *sql_variant* and *table* data types.	☑	☑	☑
Column level collations	Store objects that have different collations in the same database. Collations can be specified at the database level or at the column level.	☑	☑	☑
Integrated Analysis Services	Build OLAP, data warehousing and data mining solutions to derive value from data.	☑	☑	☑
Virtual Cube Editor	Browse, examine and edit virtual cubes graphically.	☑	☑	☑
MDX Builder	Use drag-and-drop to quickly generate Multidimensional	☑	☑	☑

[6] Active Directory is a feature of the Windows 2000 Server Family.

	Expressions. Easily add functions, arithmetic operators and more.			
Dimensions	Pick the most appropriate dimension type with new parent-child, ragged, changing, write-enabled and dependent dimensions.	☑	☑	☑
Security in Analysis Services	Lock down sensitive data by defining security for dimensions and cells. Restrict roles to specific combinations of cells.	☑	☑	☑
OLAP Actions	Allow analysis results to kick off processes in Web and business applications automatically.	☑	☑	☑
Custom rollups	Make sophisticated calculations directly from rollups with custom operators to specify how members affect parent values.	☑		